arXiv:1703.02419v1 [stat.CO] 7 Mar 2017

Probabilistic learning of nonlinear dynamical systems using sequential Monte Carlo

Thomas B. Schön[*], Andreas Svensson[†], Lawrence Murray[‡], and Fredrik Lindsten[§]

Department of Information Technology, Uppsala University

Abstract

Probabilistic modeling provides the capability to represent and manipulate *uncertainty* in data, models, decisions and predictions. We are concerned with the problem of learning probabilistic models of dynamical systems from measured data. Specifically, we consider learning of probabilistic nonlinear state space models. There is no closed-form solution available for this problem, implying that we are forced to use approximations. In this tutorial we will provide a self-contained introduction to one of the state-of-the-art methods—the particle Metropolis-Hastings algorithm—which has proven to offer very practical approximations. This is a Monte Carlo based method, where the so-called particle filter is used to guide a Markov chain Monte Carlo method through the parameter space. One of the key merits of the particle Metropolis-Hastings method is that it is guaranteed to converge to the "true solution" under mild assumptions, despite being based on a practical implementation of a particle filter (i.e., using a finite number of particles). We will also provide a motivating numerical example illustrating the method which we have implemented in an in-house developed modeling language, serving the purpose of abstracting away the underlying mathematics of the Monte Carlo approximations from the user. This modeling language will open up the power of sophisticated Monte Carlo methods, including particle Metropolis-Hastings, to a large group of users without requiring them to know all the underlying mathematical details.

1 Introduction

The true value of measured data arises once it has been analyzed and some kind of knowledge has been extracted from this analysis. The analysis often relies on the combination of a mathematical model and measured data. The model is a compact representation—set of assumptions—of some phenomenon of interest, and establishes a link between that phenomenon and the data, which is expected to provide some insight. The knowledge we seek is typically a function of some unknown variables or parameters in the model. Uncertainty plays a fundamental role in modeling since any reasonable model will be *uncertain* when making statements about unobserved variables. Probabilistic modeling allows the representation and manipulation of uncertainty in data, models, decisions and predictions. The capability to mathematically represent and manipulate uncertainty—which is essential to the development throughout this tutorial—is provided by probability theory. A good introduction to the ideas underlying probabilistic modeling in contemporary machine learning is provided by Ghahramani (2015) and from a system identification point of view we recommend Peterka (1981).

Throughout this tutorial we are concerned with the problem of learning probabilistic models of nonlinear dynamical systems from measured data, which is sometimes referred to as the nonlinear system identification problem. These models provide an interpretable representation from which it is directly possible to extract the knowledge we seek, compared to doing so directly from the data. The model is thus a natural middle ground between the measured data

[*]thomas.schon@it.uu.se
[†]andreas.svensson@it.uu.se
[‡]lawrence.murray@it.uu.se
[§]fredrik.lindsten@it.uu.se

and that knowledge. More specifically we are concerned with the nonlinear state space model

$$x_{t+1} = f(x_t, u_t, v_t, \theta), \tag{1a}$$
$$y_t = g(x_t, u_t, \theta) + e_t, \tag{1b}$$
$$\theta \sim p(\theta), \tag{1c}$$

where $y_t \in \mathsf{Y}$ denotes the observed output and u_t denotes a known input signal. In the interest of concise notation we will, without loss of generality, suppress the input signal u_t. The unknown variables are given by the state $x_t \in \mathsf{X}$ describing the system's evolution over time and the static parameters θ. The initial state is modeled as $x_0 \sim p(x_0 \mid \theta)$, where $\sim$ is used to denote the fact that x_0 is "distributed according to" some probability density function $p(x_0 \mid \theta)$. The uncertain variables v_t and e_t denote noise terms, commonly referred to as process noise and the measurement noise, respectively. Finally, the functions f and g denotes the dynamics and the measurement equation, respectively.

The first step towards extracting knowledge from a set of measured data $y_{1:T} = \{y_1, \ldots, y_T\}$ is to learn a probabilistic model of the form (1) by computing the conditional distribution of the unknown $x_{0:T}$ and θ conditioned on $y_{1:T}$, denoted $p(x_{0:T}, \theta \mid y_{1:T})$. This provides a useful representation which is typically much closer to the knowledge we seek than the measured data itself. Once we have a representation of this conditional distribution, it can be used to compute more specific quantities that typically constitute the end result of the analysis. To mention just a few examples of such quantities we have mean values, variances or estimates of some tail probability.

The key challenge is that, in general, there is no closed form expression available for $p(x_{0:T}, \theta \mid y_{1:T})$, so that we must resort to approximations. We will focus on approximations based on Monte Carlo sampling which, while being costly to compute, have the appealing property of converging to the true solution as the amount of computation increases. Over the last decade, these approximations have evolved rapidly, so that we now have computationally feasible solutions available.

This naturally brings us to the aim of this tutorial, which is to provide a gentle introduction to probabilistic learning of nonlinear dynamical systems, and to introduce in some detail one of the current state-of-the-art methods to do so. This method relies on the systematic combination of two Monte Carlo algorithms, where a sequential Monte Carlo algorithm is used to compute a good proposal distribution for a Markov chain Monte Carlo algorithm. Hence, in terms of methods, this tutorial is focused on introducing one particular solution rather than surveying all available methods (see e.g. Schön et al. (2015); Kantas et al. (2015) for recent accounts of that sort). However, the key ideas discussed in this tutorial (in the context of the specific method) are in fact central to many other state-of-the-art Monte Carlo learning methods as well. For example the particle filter itself and the fact that it is capable of producing an unbiased estimate of the likelihood are indeed of more general interest.

In the accompanying paper Svensson et al. (2017), we show how probabilistic modeling, implemented via a new algorithm of the type outlined in this tutorial, is used to solve one of the challenging benchmark problems in this issue with promising results. We will also hint at the tailored software that is being developed to make these mathematical tools available to a much wider audience without a thorough knowledge of Monte Carlo methods. Such software allows the user to focus entirely on the modeling problem and leave the computational learning problem to the software.

When we talk about a model we mean a probability distribution; in Section 2 we introduce probabilistic modeling of nonlinear state space models in more detail. The model as it is stated in (1) is clearly uncertain due to the presence of the noise sources v_t and e_t, as well as uncertainty in the initial state x_0 and in the parameters θ. In probability theory, uncertainty is represented using random variables; in Section 2, the probabilistic nature of the model will be made even more explicit when we represent it as a joint distribution $p(x_{0:T}, \theta, y_{1:T})$ of all the random variables. The basic Monte Carlo idea is then introduced in Section 3 together with an explanation of how this idea can be used to learn the conditional distribution of the parameters given the measurements, $p(\theta \mid y_{1:T})$. The idea is developed further in Section 4, where it becomes clear that we also need information about the unknown state variables, resulting in the introduction of the sequential Monte Carlo method (a.k.a. the particle filter) to estimate the state variables. The particle filter is in Section 5 used inside the MCMC method introduced in Section 3. The basic particle filter construction from Section 4 can be improved in several ways and in Section 6 we discuss some of the most important developments in this direction. The resulting method is then illustrated using a nonlinear spring-damper system in Section 7. Finally we conclude with a discussion in Section 8.

2 Probabilistic modelling of dynamical systems

We will refer to the joint distribution of all observed (here $y_{1:T}$) and unobserved (here $x_{1:T}$ and θ) variables as the *full probabilistic model*, which in our present setting amounts to $p(x_{0:T}, \theta, y_{1:T})$. The idea of using the mathematics of probability to represent and manipulate uncertainty is commonly referred to as Bayesian statistics Gelman et al. (2013). In order to write down the full probabilistic model for (1) let us start by noticing that for many models we can express the conditional distribution of y_t given x_t and θ as

$$p(y_t \mid x_t, \theta) = p_{e_t}(y_t - g(x_t, \theta), \theta), \tag{2}$$

where $p_{e_t}(\cdot)$ denotes the distribution of the measurement noise e_t. In a similar way, the dynamical equation (1a) implicitly defines a conditional probability distribution $p(x_{t+1} \mid x_t, \theta)$ describing the state evolution.[1] Hence, the SSM in (1) can be expressed as

$$x_{t+1} \mid x_t, \theta \sim p(x_{t+1} \mid x_t, \theta), \tag{3a}$$
$$y_t \mid x_t, \theta \sim p(y_t \mid x_t, \theta). \tag{3b}$$

Using conditional probability we can factor the full probabilistic model as

$$p(x_{0:T}, \theta, y_{1:T}) = \underbrace{p(y_{1:T} \mid x_{0:T}, \theta)}_{\text{data distribution}} \underbrace{p(x_{0:T} \mid \theta) p(\theta)}_{\text{prior distribution}}, \tag{4}$$

where $p(y_{1:T} \mid x_{0:T}, \theta)$ describes the distribution of the data; we refer to it as the *data distribution*. The so-called *prior distribution* $p(x_{0:T}, \theta) = p(x_{0:T} \mid \theta)p(\theta)$ represents our initial assumptions about the unknown states and parameters. It is instructive to rewrite the model (4) slightly in order to more clearly see that it is just a different way of representing (1). Let us first continue the use of conditional probabilities in order to further decompose the data distribution into

$$\begin{aligned} p(y_{1:T} \mid x_{0:T}, \theta) &= p(y_T \mid y_{1:T-1}, x_{0:T}, \theta) p(y_{1:T-1} \mid x_{0:T}, \theta) \\ &= p(y_T \mid x_T, \theta) p(y_{1:T-1} \mid x_{0:T-1}, \theta), \end{aligned} \tag{5}$$

where we made also use of the conditional independence of the observations given the current state. Moreover we can slightly rewrite the prior distribution over the states $p(x_{0:T} \mid \theta)$ by noting that

$$p(x_{0:T} \mid \theta) = p(x_T \mid x_{0:T-1}, \theta) p(x_{0:T-1} \mid \theta) = p(x_T \mid x_{T-1}, \theta) p(x_{0:T-1} \mid \theta), \tag{6}$$

where we made use of conditional probabilities in the first equality and the Markov property in the second equality. Repeated use of (5) and (6) results in

$$p(x_{0:T}, \theta, y_{1:T}) = \underbrace{\left(\prod_{t=1}^{T} \underbrace{p(y_t \mid x_t, \theta)}_{\text{observation}} \right)}_{\text{data distribution}} \underbrace{\left(\prod_{t=0}^{T-1} \underbrace{p(x_{t+1} \mid x_t, \theta)}_{\text{dynamics}} \right) \underbrace{p(x_0 \mid \theta)}_{\text{initial}} \underbrace{p(\theta)}_{\text{initial}}}_{\text{prior}} \tag{7}$$

explicitly showing how the "engineering standard" SSM, as it is formulated in (1), relates to the full probabilistic model formulation.

Starting from our model (7) (or (1)) the problem we set out to solve is how to compute the distribution of the unobserved variables conditioned on the observed variables $p(x_{0:T}, \theta \mid y_{1:T})$. This distribution is referred to as the *posterior* distribution. Using conditional probability it separates into the so-called state and parameter inference problems according to

$$p(x_{0:T}, \theta \mid y_{1:T}) = \underbrace{p(x_{0:T} \mid \theta, y_{1:T})}_{\text{state inf.}} \underbrace{p(\theta \mid y_{1:T})}_{\text{param. inf.}}. \tag{8}$$

[1] The attentive reader might wonder why the noise is assumed to be additive in (1b) but not in (1a). The reason is that the methods that we will consider require the density function $p(y_t \mid x_t, \theta)$ to be available for point-wise evaluation. The additivity assumption implies that this density can be expressed as in (2). The methods do not, however, require $p(x_{t+1} \mid x_t, \theta)$ to be available for point-wise evaluation.

Solving the parameter inference problem—which is the focus of this tutorial—involves the computation of

$$p(\theta \mid y_{1:T}) = \frac{p(y_{1:T} \mid \theta)p(\theta)}{p(y_{1:T})} = \frac{p(y_{1:T} \mid \theta)p(\theta)}{\int p(y_{1:T} \mid \theta)p(\theta)\mathrm{d}\theta}. \tag{9}$$

The distribution that we are aiming to find is referred to as the *target* distribution and it is denoted by $\pi(\cdot)$. The target distribution throughout this tutorial will mainly be $\pi(\theta) = p(\theta \mid y_{1:T})$. The end user may not be interested in this target distribution *per se*, but rather in some test function $\varphi(\theta)$ evaluated over it. These can be evaluated by solving integrals of the form

$$\mathbb{E}\left[\varphi(\theta) \mid y_{1:T}\right] = \int \varphi(\theta)p(\theta \mid y_{1:T})\mathrm{d}\theta. \tag{10}$$

A common test function is just the identity function $\varphi(\theta) = \theta$, so that the integral (10) provides an estimate of the expected value of θ conditioned on the data $y_{1:T}$, i.e. $\mathbb{E}\left[\theta \mid y_{1:T}\right]$. Another example is the indicator function $\varphi(\theta) = I(\theta > \vartheta)$, for some threshold value ϑ, which provides an estimate of a tail probability, perhaps important in modelling extreme events. Other test functions or combinations of them yield the covariance and higher moments, or estimates of domain-specific utility or loss.

From (9), it is clear that in order to compute the posterior distribution of the unknown parameters we first need to compute the data distribution $p(y_{1:T} \mid \theta)$. It is commonly computed by averaging (marginalizing) the joint distribution $p(x_{0:T}, y_{1:T} \mid \theta)$ over all possible values for the state variables $x_{0:T}$

$$p(y_{1:T} \mid \theta) = \int p(y_{1:T}, x_{0:T} \mid \theta)\mathrm{d}x_{0:T}. \tag{11}$$

Besides being useful to us later, this indicates the important fact that the state inference problem in inherent in the parameter inference problem when we are dealing with SSMs. Note that if we instead model the unknown parameters θ as deterministic variables to be estimated using maximum likelihood, we still have to deal with the high-dimensional integral in (11), since this will now define the likelihood. The *likelihood* is defined as the data distribution evaluated for the particular measurement sequence $y_{1:T}$ we have available. Hence, the computation of maximum likelihood estimates will benefit a lot from the methodology presented in this paper as well.

We make use of $p(\cdot)$ to denote a probability density function of variables that in some way describes the model or can be induced from the model (1).

Distributions describing random variables that are not directly related to the model are instead denoted using characters from the greek alphabet. These variables will for example arise as parts of the (randomized) algorithms that we will derive in the coming sections.

3 Solving the parameter inference problem

The parameter inference problem amounts to computing the posterior distribution $p(\theta \mid y_{1:T})$ or more commonly the problem is to evaluate the integral (10) with respect to this distribution for some test function $\varphi(\theta)$. This integral does not allow for closed-form solutions in general[2]. We can, however, construct an estimator of it using the Monte Carlo idea, which is introduced in Section 3.1. More specifically we will make use of a so-called pseudo-marginal Markov chain Monte Carlo (MCMC) method— introduced in Section 3.3—which itself employs an estimate. However, before we describe the pseudo-marginal idea we will first introduce the basic MCMC idea itself in Section 3.2.

[2]For certain, indeed important, special cases such as linear models with only Gaussian noise, closed-form solutions exists.

3.1 The Monte Carlo idea

Monte Carlo integration provides approximate solutions to integrals of the form

$$c = \mathbb{E}\left[h(\xi)\right] \triangleq \int_{\Xi} h(\xi) p_{\xi}(\xi) \mathrm{d}\xi, \tag{12}$$

where $\xi \in \Xi$ and $\xi \sim p_{\xi}(\xi)$. A classical choice to evaluate the integral is to take N equally-spaced grid points $\{\xi^n\}_{n=1}^N$ on some interval $[a, b]$, and compute

$$\widehat{c} = \frac{1}{N} \sum_{n=1}^{N} h(\xi^n) p(\xi^n) \tag{13}$$

as an estimate of c. This is the usual Riemann sum. In Monte Carlo integration, the idea is to instead choose *random* sample points $\{\xi^n\}_{n=1}^N$ from the probability distribution $p_{\xi}(\xi)$ and compute

$$\widehat{c} = \frac{1}{N} \sum_{n=1}^{N} h(\xi^n). \tag{14}$$

Clearly $\widehat{c}$ is now itself random. There are two advantages to this approach. The first is that, while a fixed grid can give an arbitrarily bad estimate, it can be proven that the expectation of the random estimator $\widehat{c}$ from Monte Carlo integration is exactly c, i.e., $\mathbb{E}\left[\widehat{c}\right] = c$, and we say that $\widehat{c}$ is an *unbiased* estimator of c when this hold true. Secondly, if the dimension of Ξ is D, the error of the Riemann sum (13) scales as $\mathcal{O}(N^{-1/D})$, while that of the Monte Carlo estimate (14) scales as $\mathcal{O}(N^{-1/2})$ (see e.g. Robert and Casella, 2004; Owen, 2013). That is, for more than a few dimensions, the Monte Carlo estimator exhibits smaller error.

3.2 The Markov chain Monte Carlo idea

A Markov chain $\{\theta[m]\}_{m\geq 0}$ is a stochastic process with the property that the next state $\theta[m+1]$ of the process depends on the current state $\theta[m]$, but it is conditionally independent of all previous states $\theta[0:m-1]$. It is completely specified by an initial distribution and a transition kernel. A stationary Markov chain converge to a particular *stationary distribution*. The SSM (1) is a popular and useful example of a Markov chain $\{x_t\}_{t\geq 0}$ with initial distribution $p(x_0)$ and transition kernel given by $p(x_{t+1} \mid x_t)$.

The idea underlying MCMC is to simulate a Markov chain which is designed in such a way that its stationary distribution coinsides with the target distribution, here taken to be $\pi(\theta) = p(\theta \mid y_{1:T})$. Our interest lies in Markov chains which are structured such that as $M \to \infty$ we have that

$$\frac{1}{M} \sum_{m=0}^{M} \varphi(\theta[m]) \xrightarrow{\text{a.s.}} \int \varphi(\theta) p(\theta \mid y_{1:T}) \mathrm{d}\theta, \tag{15}$$

where $\xrightarrow{\text{a.s.}}$ denotes almost sure convergence. The finite and explicit sum in (15) does indeed provide an approximation of the intractable integrals *if* there is a way of constructing the underlying Markov chain $\{\theta[m]\}_{m\geq 0}$. The construction of such a Markov chain is possible and it consists of two steps. In the first step a new *candidate* state θ' is proposed by generating a sample from the so-called *proposal distribution* $q(\theta \mid \theta[m])$, i.e. $\theta' \sim q(\theta \mid \theta[m])$. The proposal distribution depends on the current state $\theta[m]$ and its exact form is decided by the user, with the only requirement that the support of the proposal distribution contains the support of the target distribution, i.e. $q(\theta \mid \theta[m]) > 0$ whenever $\pi(\theta) > 0$.

In the second step we have to decide if the next state $\theta[m+1]$ of the Markov chain should be the candidate sample θ' just produced or if it is better to stay in the current state $\theta[m]$. It can be shown that if we chose the candidate sample as the next state with probability

$$\alpha = \min\left(1, \frac{\pi(\theta')}{\pi(\theta[m])} \frac{q(\theta[m] \mid \theta')}{q(\theta' \mid \theta[m])}\right) = \min\left(1, \frac{p(y_{1:T} \mid \theta') p(\theta')}{p(y_{1:T} \mid \theta[m]) p(\theta[m])} \frac{q(\theta[m] \mid \theta')}{q(\theta' \mid \theta[m])}\right), \tag{16}$$

the resulting Markov chain will have our target distribution as its stationary distribution. The second equality in (16) is due to (9). For natural reasons α in (16) is referred to as the *acceptance probability*. Indeed, with probability α the new state in the Markov chain is chosen as $\theta[m+1] = \theta'$ (the candidate sample is accepted) and with probability $1-\alpha$ the candidate sample is rejected and the Markov chain remains in the same state as in the previous iteration, i.e. $\theta[m+1] = \theta[m]$. The samples $\{\theta[m]\}_{m=k}^{M}$ from the resulting Markov chain constitutes an empirical approximation

$$\widehat{\pi}(\theta) = \frac{1}{M}\sum_{m=0}^{M} \delta_{\theta[m]}(\theta) \tag{17}$$

of the posterior distribution. Here, $\delta_{\theta[m]}(\theta)$ denotes the (Dirac) point-mass distribution at $\theta = \theta[m]$. The algorithm of first proposing a candidate state and then accepting or rejecting this as the next state of the Markov chain is referred to as the *Metropolis-Hastings (MH) algorithm* and it was introduced by Metropolis et al. (1953); Hastings (1970) and there are by now many textbook style introductions available, see e.g. Robert and Casella (2004); Owen (2013).

The intractable integral (10) is now turned into a tractable finite sum simply by inserting (17) into (10), resulting in

$$\widehat{\pi}_\theta[\varphi] \triangleq \int \varphi(\theta)\widehat{\pi}(\theta)\mathrm{d}\theta = \frac{1}{M}\sum_{m=0}^{M}\int \varphi(\theta)\delta_{\theta[m]}(\theta)\mathrm{d}\theta = \frac{1}{M}\sum_{m=0}^{M}\varphi(\theta[m]), \tag{18}$$

where the last equality follows from the properties of the Dirac function. The estimator $\widehat{\pi}_\theta[\varphi]$ defined above is well-behaved in the sense that it obeys both a law of large numbers and a central limit theorem, see e.g. Robert and Casella (2004); Meyn and Tweedie (2009) for a thorough treatment. The law of large numbers tells us that the estimator is consistent and asymptotically unbiased as $M \to \infty$, whereas the central limit theorem tells us that the error is approximately Gaussian with a variance that decrease with the standard Monte Carlo rate of $1/M$. A nice historical account of MCMC is provided in Robert and Casella (2011).

3.3 Using unbiased estimates within Metropolis-Hastings

The problem preventing us from implementing the Metropolis-Hastings algorithm for a general state space model is that we cannot compute the acceptance probability α given in (16), since there is no closed-form expression available for the likelihood $p(y_{1:T} \mid \theta)$. However, what if we have an estimate of the likelihood, can we then use this estimate in place of the exact likelihood in the computation of the acceptance probability and still end up with a valid algorithm? ("Valid", here, refers to that the method converges in the sense of (15).) The answer to this highly nontrivial question is actually *yes*—under certain conditions on the likelihood estimate (described below). The result is referred to as an *exact approximation*, since we obtain an *exact* Metropolis-Hastings algorithm—in the sense that the target distribution of interest remains the stationary distribution of the Markov chain—despite the fact that we make use of an *approximation* of the likelihood when evaluating the acceptance probability.

Let us now provide some intuition as to why this actually works. We start by assuming that we have access to an estimator z of the likelihood $p(y_{1:T} \mid \theta)$. The estimator naturally depends on the observed data $y_{1:T}$ and the model parameter θ (since these variables determine the value of the estimand). Furthermore, we assume that the estimator depends on some additional random variables—as we shall see below, these random variables typically come from yet another Monte Carlo procedure which is used to compute the estimate. Consequently, the estimator z is itself a *random variable* and it has some distribution $\psi(z \mid \theta, y_{1:T})$ depending on θ and $y_{1:T}$. (Note that we will not require this distribution to be available on closed form, but it exists conceptually nevertheless.)

Next, we introduce the random variable z (i.e., the estimator of the data distribution) as an *auxiliary variable* in the model. Auxiliary variables—while being of no particular interest on their own—are commonly introduced in statistical models in order to simplify the inference for some other variable of interest. In our case, the variable of interest is the model parameter θ. Thus, we consider the joint distribution for (θ, z) given by

$$\psi(\theta, z \mid y_{1:T}) = p(\theta \mid y_{1:T})\psi(z \mid \theta, y_{1:T}) = \frac{p(y_{1:T} \mid \theta)p(\theta)\psi(z \mid \theta, y_{1:T})}{p(y_{1:T})}. \tag{19}$$

We note that the original target distribution $p(\theta \mid y_{1:T})$ is (by construction) obtained by marginalizing the joint distribution $\psi(\theta, z \mid y_{1:T})$ with respect to the auxiliary variable z.

Algorithm 1 pseudo-marginal Metropolis-Hastings

1: **Initialisation ($m = 0$):** Set $\theta[0]$ arbitrary.
2: **for** $m = 1$ **to** M **do**
3: Sample $\theta' \sim q(\theta \,|\, \theta[m-1])$.
4: Sample $z' \sim \psi(z \,|\, \theta', y_{1:T})$.
5: With probability

$$\alpha = \min\left(1, \frac{z' p(\theta')}{z[m-1]p(\theta[m-1])} \frac{q(\theta[m-1] \,|\, \theta')}{q(\theta' \,|\, \theta[m-1])}\right), \tag{23}$$

set $\{\theta[m], z[m]\} \leftarrow \{\theta', z'\}$ (accept the candidate samples) and with probability $1-\alpha$ set $\{\theta[m], z[m]\} \leftarrow \{\theta[m-1], z[m-1]\}$ (reject the candidate samples).
6: **end**

If we now were to construct a Metropolis-Hastings algorithm for both the parameters and the auxiliary variable it would suffer from the same problem as before, since the (intractable) likelihood still appears in the expression (19). To overcome this difficulty we will make use of the following trick: we *define* a new joint distribution over (θ, z) by simply replacing the intractable likelihood $p(y_{1:T} \,|\, \theta)$ in (19) with its estimator z

$$\phi(\theta, z \,|\, y_{1:T}) := \frac{z\, p(\theta)\psi(z \,|\, \theta, y_{1:T})}{p(y_{1:T})}. \tag{20}$$

For this to be a valid probability distribution we have two formal requirements: *(i)* that it is everywhere non-negative, and *(ii)* that it integrates to 1. Furthermore, if this distribution is to be of any use to us, it also has to preserve the property that the marginal distribution for θ with respect to ϕ coincides with the target distribution of interest: $\int \phi(\theta, z \,|\, y_{1:T}) dz = p(\theta \,|\, y_{1:T})$. As we shall see next, all of these requirements are fulfilled if z is a *non-negative* and *unbiased* estimate of $p(y_{1:T} \,|\, \theta)$.

The first requirement—non-negativity of ϕ—follows directly from the assumed non-negativity of z. For the second and third requirements, we consider the integral

$$\int \phi(\theta, z \,|\, y_{1:T}) \mathrm{d}z = \frac{p(\theta)}{p(y_{1:T})} \int z\, \psi(z \,|\, \theta, y_{1:t}) \mathrm{d}z. \tag{21}$$

However, since $\psi(z \,|\, \theta, y_{1:T})$ is nothing but the distribution of the estimator z, the integral in the expression above is just the mean of z. Hence, due to the assumed unbiasedness of z

$$\int \phi(\theta, z \,|\, y_{1:T}) \mathrm{d}z = \frac{p(\theta)}{p(y_{1:T})} p(y_{1:T} \,|\, \theta) = p(\theta \,|\, y_{1:T}). \tag{22}$$

Thus, we see that the marginal distribution of θ with respect to ϕ is $p(\theta \,|\, y_{1:T})$ as required, which also implies that ϕ is a properly normalised probability distribution (it integrates to 1, which can be seen by further integrating both sides of (22) with respect to θ).

The idea of making use of an unbiased estimate of the likelihood within a Metropolis-Hastings algorithm is referred to as the pseudo-marginal approach and the resulting algorithm as it was explained above is summarized in Algorithm 1. It was first introduced in 2003 Beaumont (2003). It has later been generalized and thoroughly analysed, see Andrieu and Roberts (2009); Andrieu and Vihola (2015).

Note that the pseudo-marginal algorithm will converge towards the correct posterior distribution $p(\theta \,|\, y_{1:T})$ under weak assumptions, as long as the nonnegativity and unbiasedness conditions hold for the estimate z (convergence, here, means that the distribution of $\theta[m]$ converges to $p(\theta \,|\, y_{1:T})$ as $m \to \infty$). This is not to say that the precision of the estimate z is of no interest, though. Indeed, if the variance in the estimate z is overly large, then this will result in a slow convergence speed of the algorithm. In practice it is therefore important to keep the variance in the estimate as low as possible to obtain an efficient solution. We return to this in Section 5.

4 Computing an unbiased estimate of the likelihood using the particle filter

In order to use Algorithm 1 for identification of a state space model, we need a way to construct an *unbiased* and *non-negative* estimate of the likelihood $p(y_{1:T} \mid \theta)$ for any value of θ. Specifically, as the algorithm considers a candidate value θ' for the model parameters, we need to compute an estimate $z \approx p(y_{1:T} \mid \theta')$ (lines 3–4) that is used to determine whether or not θ' is a promising candidate.

The likelihood is only a function of the data $y_{1:T}$ and the parameters θ, and not the state variables $x_{0:T}$. Thus, to compute $p(y_{1:T} \mid \theta)$ we need to marginalize—average—out all possible trajectories $x_{0:T}$. We can write this as

$$p(y_{1:T} \mid \theta) = \int_{\mathsf{X}^{T+1}} p(y_{1:T} \mid x_{0:T}, \theta) p(x_{0:T} \mid \theta) \mathrm{d}x_{0:T}. \tag{24}$$

In general, this integral admits no closed-form solution[3]. We can, however, construct an unbiased and non-negative estimator of it, yet again[4] using the Monte Carlo idea. We will in Section 4.1 detail a naive construction of such an estimator. That construction serves as a motivation for the particle filter, which is then introduced in Section 4.2. It provides a practically useful computation of a non-negative and unbiased estimator of the likelihood for a nonlinear state space model (1).

4.1 Monte Carlo integration for the data distribution

The problem of solving (24) is a special case of (12) where $\xi = x_{0:T}$, $p_\xi(\xi) = p(x_{0:T} \mid \theta)$ and $h(\xi) = p(y_{1:T} \mid x_{0:T}, \theta)$. A first attempt to solve this problem is to use the basic Monte Carlo approach as presented in Section 3.1 and proceed as follows:

1. Generate N samples of the state trajectory $x_{0:T}^n \sim p(x_{0:T} \mid \theta)$ for $n = 1, \ldots, N$. In practice, this can be done by simulating the system dynamics: $x_0^n \sim p(x_0 \mid \theta)$, then $x_1^n \sim p(x_1 \mid x_0^n, \theta)$, then $x_2^n \sim p(x_2 \mid x_1^n, \theta)$, etc. for $n = 1, \ldots, N$.
2. Compute an estimate of the likelihood as $z := \frac{1}{N} \sum_{n=1}^N p(y_{1:T} \mid x_{0:T}^n, \theta) = \frac{1}{N} \sum_{n=1}^N \prod_{t=1}^T p(y_t \mid x_t^n, \theta)$.

This results in a basic Monte Carlo estimate z which is indeed unbiased and non-negative so it is a valid approach. In practice, however, the variance of this z is too large to be of any practical use, as we will illustrate in Figure 1. The reason is that many samples $\{x_{0:T}^n\}_{n=1}^N$ are likely to be drawn in a place where they contribute very little to the estimate due to the high dimension of the space X^T.

To mitigate this problem the structure of the state space model can be leveraged. Specifically, we factorise the high-dimensional integral into a product of T lower-dimensional integrals:

$$p(y_{1:T} \mid \theta) = \prod_{t=1}^T p(y_t \mid y_{1:t-1}, \theta) = \prod_{t=1}^T \int_{\mathsf{X}} p(y_t \mid x_t, \theta) p(x_t \mid y_{1:t-1}, \theta) \mathrm{d}x_t. \tag{25}$$

We can then perform Monte Carlo integration for each integral, one at a time, and take the product of the estimates to obtain an overall estimate. The advantage here is that it is typically much easier to handle these lower-dimensional integrals than it is to handle the preceding higher-dimensional integral. The drawback, however, is that it is typically not possible to simulate directly from $p(x_t \mid y_{1:t-1}, \theta)$. To tackle this issue we will instead make use of an algorithm known as the *particle filter*, which we introduce next.

[3] For certain, indeed important, special cases such as linear models with only Gaussian noise, closed-form solutions exists, obtained via the Kalman filter. Complete details on that special case—including how to use Kalman filters to implement a Metropolis-Hastings algorithm—is provided in Schön et al. (2015).

[4] The first time was the Metropolis-Hastings algorithm in the previous section.

4.2 Introducing the particle filter

The particle filter sequentially generates samples $\{x_t^n\}_{n=1}^N$ for $t = 1, 2, \ldots$, such that the t^{th} set of samples are approximately distributed according to $p(x_t \mid y_{1:t-1}, \theta)$. What differs from the basic Monte Carlo procedure outlined above, however, is that all the samples are allowed to interact between iterations. To derive the particle filter it is convenient to introduce the notion of an *empirical approximation* of a probability distribution. Based on the N samples $\{x_t^n\}_{n=1}^N$ we can approximate the distribution $p(x_t \mid y_{1:t-1}, \theta)$ as

$$\widehat{p}(x_t \mid y_{1:t-1}, \theta) = \frac{1}{N} \sum_{n=1}^{N} \delta_{x_t^n}(x_t), \tag{26}$$

where $\delta_{x_t^n}(x_t)$ is again a (Dirac) point-mass distribution at x_t^n. Note that if we plug this empirical distribution into an integral with respect to $p(x_t \mid y_{1:t-1}, \theta)$ we recover the normal Monte Carlo estimator of that integral. For instance, for the integral in (25) we get,

$$\int_{\mathsf{X}} p(y_t \mid x_t, \theta) p(x_t \mid y_{1:t-1}, \theta) \mathrm{d}x_t \approx \int_{\mathsf{X}} p(y_t \mid x_t, \theta) \left[\frac{1}{N} \sum_{n=1}^{N} \delta_{x_t^n}(x_t) \right] \mathrm{d}x_t = \frac{1}{N} \sum_{n=1}^{N} p(y_t \mid x_t^n, \theta). \tag{27}$$

Let us now consider how the samples—or particles as they are often called—$\{x_t^n\}_{n=1}^N$ are generated sequentially. When $t = 1$ we have, by using the convention $y_{1:0} = \emptyset$, that $p(x_1 \mid y_{1:0}, \theta) = p(x_1 \mid \theta)$ and we can thus sample directly from the initial distribution $x_0^n \sim p(x_0 \mid \theta)$ and simulate the system dynamics to obtain $x_1^n \sim p(x_1 \mid x_0^n)$ for $n = 1, \ldots, N$. For any consecutive time step it holds that each target distribution $p(x_{t+1} \mid y_{1:t}, \theta)$ can be constructed from the previous, $p(x_t \mid y_{1:t-1}, \theta)$, in the following way. First, by Bayes' theorem it follows that

$$p(x_t \mid y_{1:t}, \theta) \propto p(y_t \mid x_t, \theta) p(x_t \mid y_{1:t-1}, \theta). \tag{28}$$

Next, by incorporating the state at time $t + 1$ (according to the dynamical model) and marginalising over x_t we have

$$p(x_{t+1} \mid y_{1:t}, \theta) = \int p(x_{t+1}, x_t \mid y_{1:t}, \theta) \mathrm{d}x_t = \int p(x_{t+1} \mid x_t, \theta) p(x_t \mid y_{1:t}, \theta) \mathrm{d}x_t. \tag{29}$$

The two equations above are often referred to as the measurement update and time update, respectively, as the former takes the t^{th} measurement y_t into account and the latter propagates the distribution forward in time according to the system dynamics (cf. the two steps of the Kalman filter).

The particle filter makes use of Equations (28) and (29) to sample approximately from $p(x_{t+1} \mid y_{1:t}, \theta)$ based on the existing samples from $p(x_t \mid y_{1:t-1}, \theta)$. First, we plug the empirical approximation (26) into (28). This gives rise to an empirical approximation of the distribution $p(x_t \mid y_{1:t}, \theta)$, but where each particle is assigned a *weight* w_t^n according to the multiplicative factor $w_t^n = p(y_t \mid x_t^n, \theta)$. The unknown normalising constant in (28) (hidden in the proportionality sign) is not needed since we can simply normalise the weights of the empirical distribution, which thus becomes

$$\widehat{p}(x_t \mid y_{1:t}, \theta) = \sum_{n=1}^{N} \frac{w_t^n}{\sum_{j=1}^{N} w_t^j} \delta_{x_t^n}(x_t). \tag{30}$$

Next, we note that by (29) it is *conceptually* possible to sample from $p(x_{t+1} \mid y_{1:t}, \theta)$ by first sampling from $p(x_t \mid y_{1:t}, \theta)$ and then simulating the system dynamics forward by sampling from the transition density $p(x_{t+1} \mid x_t, \theta)$. To make this practical, the particle filter plugs the empirical approximation (30) into (29), resulting in a weighted empirical approximation. Thus, we simulate N particles from the empirical distribution (30). This simply amounts to sampling with replacement from among the existing particles $\{x_t^n\}_{n=1}^N$, with probabilities given by the normalised weights. Let the resulting particles be denoted by $\{\bar{x}_t^n\}_{n=1}^N$. Then we propagate these particle forward in time by sampling $x_{t+1}^n \sim p(x_{t+1}^n \mid \bar{x}_t^n, \theta)$ for $n = 1, \ldots, N$. This completes one time step of the particle filter, and the procedure described above can now be repeated for time $t + 2$, $t + 3$, etc.

The procedure of sampling from the empirical distribution (30) is commonly referred to as *resampling*. This is a key ingredient of the particle filter which, intuitively, exchanges the weight of each particle for a number of copies among

Algorithm 2 Bootstrap particle filter (all operations are for $n = 1, \ldots, N$)

1: **Initialisation:**
2: Sample $x_0^n \sim p(x_0 \mid \theta)$ and propagate $x_1^n \sim p(x_1 \mid x_0^n, \theta)$.
3: Compute $w_1^n = p(y_1 \mid x_1^n, \theta)$.
4: **for** $t = 2$ **to** T **do**
5: **Resampling:** Sample a_t^n with $\mathbb{P}(a_t^n = j) \propto w_{t-1}^j$ and set $\bar{x}_{t-1}^n = x_{t-1}^{a_t^n}$.
6: **Propagation:** Sample $x_t^n \sim p(x_t \mid \bar{x}_{t-1}^n, \theta)$.
7: **Weighting:** Compute $w_t^n = p(y_t \mid x_t^n, \theta)$.
8: **end**

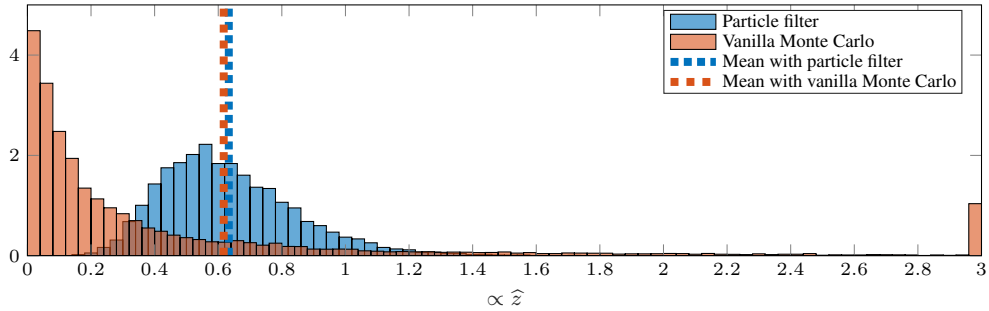

Figure 1: Estimation of the likelihood z in the spring-damper example (Section 7). The histograms describes how 10 000 samples obtained by vanilla Monte Carlo integration (orange, as described in Section 4.1) and the particle filter (blue, as described in Section 4.2) were distributed. Both approaches provides an unbiased estimate (the means, dotted lines, are indeed essentially the same), but the particle filter has significantly smaller variance and less heavy tails than the vanilla Monte Carlo integration (the rightmost orange bin contains all samples ≥ 3, and the biggest orange sample obtained was as large as 100.)

the N chosen. Particles with relatively small weights tend to be discarded, while particles with relatively large weights tend to be replicated several times.

We summarize this by Algorithm 2, the *bootstrap particle filter*, which was introduced independently by Gordon et al. (1993); Stewart and McCarty (1992); Kitagawa (1996).

From each empirical approximation $\widehat{p}(x_t \mid y_{1:t-1}, \theta)$, it is straightforward to perform Monte Carlo integration of (25) by substituting the empirical approximation $\widehat{p}(x_t \mid y_{1:t-1}, \theta)$ in place of each $p(x_t \mid y_{1:t-1}, \theta)$ as in (27). This results in an estimate of the likelihood given by

$$z := \prod_{t=1}^{T} \left[\frac{1}{N} \sum_{n=1}^{N} p(y_t \mid x_t^n, \theta) \right] = \prod_{t=1}^{T} \left[\frac{1}{N} \sum_{n=1}^{N} w_t^n \right]. \tag{31}$$

(Note that the weights w_t^n appearing in the above expression are not normalised.) This estimator is obviously non-negative, since each term w_t^n is non-negative. What is less obvious to see is that the estimator also satisfies the requirement of being unbiased, which means that it can be used within Algorithm 1 to obtain a valid pseudo-marginal Metropolis-Hastings method. For a proof of this claim we refer the interested reader to Del Moral (2004); Pitt et al. (2012).

We summarize this section by Figure 1, where we have implemented the approach proposed in the beginning of Section 4.1 ('vanilla Monte Carlo'), as well as the bootstrap particle filter, to illustrate the different properties of the likelihood estimate z for both approaches. As clearly can be seen, the variance of z obtained from the particle filter is much smaller than the variance obtained by the vanilla Monte Carlo approach. Nevertheless, both approaches provide an unbiased estimate.

5 Using the particle filter estimate within pseudo-marginal MH

The non-negative and unbiased likelihood estimator described in Section 4.2 can now be used inside the pseudo-marginal Metropolis-Hastings algorithm from Section 3.3. The end result is a standard Metropolis-Hastings algorithm operating on the non-standard joint space of the model parameters θ and the auxiliary variable z. The target distribution for our new algorithm is given by $\phi(\theta, z \mid y_{1:T})$, which was defined in (20). The result is referred to as *particle Metropolis-Hastings (PMH)*—summarized in Algorithm 3. It was introduced by Andrieu et al. in 2010 Andrieu et al. (2010) (under the name particle *marginal* Metropolis-Hastings).

Algorithm 3 Particle Metropolis-Hastings (PMH)

1: **Initialisation ($m = 0$):** Set $\theta[0]$ arbitrary and run Algorithm 2 to obtain $z[0]$.
2: **for** $m = 1$ **to** M **do**
3: Sample $\theta' \sim q(\theta \mid \theta[m-1])$.
4: Sample $z' \sim \psi(z \mid \theta', y_{1:T})$ by running Algorithm 2 once and compute (31).
5: Compute the acceptance probability

$$\alpha_m = \min\left(1, \frac{z' p(\theta')}{z[m-1]p(\theta[m-1])} \frac{q(\theta[m-1] \mid \theta')}{q(\theta' \mid \theta[m-1])}\right), \tag{32}$$

6: Sample ω_m uniformly over $[0, 1]$.
7: **if** $\omega_m < \alpha_m$ **then**
8: Set $\{\theta[m], z[m]\} \leftarrow \{\theta', z'\}$ (accept the candidate samples).
9: **else**
10: Set $\{\theta[m], z[m]\} \leftarrow \{\theta[m-1], z[m-1]\}$ (reject the candidate samples).
11: **end**

The PMH algorithm will thus produce a Markov chain $\{\theta[m], z[m]\}_{m=0}^M$ on the joint space of the parameters θ and the auxiliary variable z. Since the posterior distribution $p(\theta \mid y_{1:T})$ is obtained as a marginal of $\phi(\theta, z \mid y_{1:T})$ we can—by construction—obtain samples from $p(\theta \mid y_{1:T})$ by extracting the sub-chain $\{\theta[m]\}_{m=0}^M$ from $\{\theta[m], z[m]\}_{m=0}^M$. Hence, the samples from the auxiliary variable are simply ignored. It is worth noting that the samples of the states $x_{0:T}$ produced by the particle filters that are executed as parts of Algorithm 3 provide a competitive solution to the problem of estimating the joint smoothing distribution $p(x_{0:T} \mid y_{1:T}, \theta)$.

As pointed out in Section 3, the variance of the estimate z will affect the convergence speed of the algorithm. Specifically, if the variance is very large, then the method tends to "get stuck" for many consecutive iterations, not accepting any proposed moves. The reason for this is that large variance in the estimate is often related to a large skewness as well, in the sense that with high probability $z < p(y_{1:T} \mid \theta)$. However, since the estimate is unbiased, this implies that z sometimes (but rarely) must take on values $z \gg p(y_{1:T} \mid \theta)$. This is illustrated in Figure 1 where the vanilla Monte Carlo estimate is heavily skewed. This is problematic when used in Algorithm 3 since if $z[m-1] \gg p(y_{1:T} \mid \theta)$ in (32), the acceptance probability α_m will tend to be small and the sampler can get stuck for many iterations. For the method to work satisfactorily we therefore need to ensure that the variance in z is not overly large, which in turn implies that we need to use a sufficiently large number N of particles in the underlying particle filter. What "sufficiently" means is problem dependent, but as a rule-of-thumb N should scale linearly with T.

A possible improvement of Algorithm 3 is offered by noting that the particle filter output can also be used to compute estimates of the likelihood gradient and Hessian. These estimates can be used to construct better proposal distributions for θ, with the potential to explore the parameter space similarly to the way in which Newton's optimization algorithm is exploring the parameter space. This can lead so significant improvements compared to the standard random walk proposal that is commonly employed, see Dahlin et al. (2015). The classic variance reduction technique of positively correlating two estimators has been applied also to the particle Metropolis Hastings algorithm by introducing positive correlations between subsequent likelihood estimators Deligiannidis et al. (2015).

6 Variance reduction methods for the particle filter

As illustrated by the example above, the bootstrap particle filter results in a substantial improvement in the estimate of the likelihood over the basic Monte Carlo approach (presented in Section 4.2). Since both approaches result in unbiased estimates, this improvement essentially stems from a reduction in variance. Various improvements on the bootstrap particle filter can similarly be used to further reduce the variance of the estimate of the data distribution (while maintaining unbiasedness). Indeed, many such improvements have been proposed in the literature over the past quarter of a century. For the interested reader we review a few key ideas below. Additional improvements and details can be found in e.g. Doucet and Johansen (2011); Cappé et al. (2007); Doucet et al. (2000).

6.1 Low-variance resampling

Firstly, we note that the resampling step of the particle filter—while being of key importance to its stability—is a random procedure. As such it inevitably introduces some variance in the estimate of the likelihood. This additional variance can be reduced by using standard variance reduction techniques, e.g. stratification, when sampling from the empirical distribution (30). This still results in a valid particle filter, as long as the resampling method used is itself unbiased. More precisely, when sampling from the empirical distribution (30) we require that the expected number of copies of each particle is proportional to its weight,

$$\mathbb{E}\left[\frac{1}{N}\sum_{n=1}^{N} I(\bar{x}_t^n = x_t^i)\right] = \frac{w_t^i}{\sum_{n=1}^{N} w_t^n}, \qquad i = 1, \ldots, N, \tag{33}$$

where $I(\cdot)$ is an indicator function. Several low-variance resampling methods satisfying this unbiasedness condition are reviewed in Douc et al. (2005); Hol et al. (2006).

Low-variance resampling can (and should) always be used when implementing the particle filter. Other variance reduction techniques, however, are more model specific and can be used only for certain classes of state space models. We describe two such possible improvements below.

6.2 Conditioning on y_t

The first is a technique which aims to take the t^{th} measurement y_t into account when simulating the particles $\{x_t^n\}_{n=1}^N$ at time t. The intuition behind this is that the measurement y_t often contains valuable information about the state of the system at time t, which can help to simulate these particles in a "good region" of the state space. To this end we assume that the distribution for the system dynamics *conditionally* on the current observation, i.e.,

$$p(x_t \,|\, x_{t-1}, y_t, \theta) = \frac{p(y_t \,|\, x_t, \theta) p(x_t \,|\, x_{t-1}, \theta)}{p(y_t \,|\, x_{t-1}, \theta)}, \tag{34}$$

is available on closed form. This is not always the case, but for some highly relevant state space models this distribution is indeed available. One important example is when the state dynamics $p(x_t \,|\, x_{t-1}, \theta)$ is Gaussian (but with a possibly nonlinear dependence on x_{t-1}) and the measurement equation is linear and Gaussian, $p(y_t \,|\, x_t, \theta) = \mathcal{N}(y_t \,|\, Cx_t, R)$. We shall further assume that the normalisation factor $p(y_t \,|\, x_{t-1}, \theta)$ in (34) can be evaluated point-wise, but this typically follows from the aforementioned assumption.

The effect of simulating particles x_t^n conditionally on y_t is that we obtain an (unweighted) empirical approximation of the *filtering distribution* instead of the one-step predictive distribution as in (27). That is, we sequentially obtain the empirical approximations

$$\widehat{p}(x_t \,|\, y_{1:t}, \theta) = \frac{1}{N}\sum_{n=1}^{N} \delta_{x_t^n}(x_t), \tag{35}$$

for $t = 0, \ldots, T$. Note the conditioning on y_t on the left-hand side in (35).

Our main object of interest in this paper—the likelihood—can be expressed in terms of these filtering distributions using a similar factorisation as in (25), but where we choose to marginalise over x_{t-1} instead of x_t,

$$p(y_{1:T} \mid \theta) = \prod_{t=1}^{T} \int_{\mathsf{X}} p(y_t \mid x_{t-1}, \theta) p(x_{t-1} \mid y_{1:t-1}, \theta) \mathrm{d}x_{t-1}. \tag{36}$$

To see how we can sequentially obtain the empirical distributions (35) we start by combining the measurement and time update equations, (28) and (29), with the expression (34),

$$\begin{aligned} p(x_t \mid y_{1:t}, \theta) &\propto \int p(y_t \mid x_t) p(x_t \mid x_{t-1}, \theta) p(x_{t-1} \mid y_{1:t-1}, \theta) \mathrm{d}x_{t-1} \\ &= \int p(x_t \mid x_{t-1}, y_t, \theta) p(y_t \mid x_{t-1}, \theta) p(x_{t-1} \mid y_{1:t-1}, \theta) \mathrm{d}x_{t-1}. \end{aligned} \tag{37}$$

Proceeding in a similar way as for the bootstrap particle filter presented above, we plug in the "current" empirical approximation of the filtering distribution (i.e., from time $t-1$) into the above integral. In doing so, the factor $p(y_t \mid x_{t-1}, \theta)$ enters as a weight on the particles at time $t-1$. Thus, to sample a new set of particles $\{x_t^n\}_{n=1}^N$ approximately from the filtering distribution at time t, we *(i)* resample the particles at time $t-1$ with probabilities proportional to $\nu_t^n := p(y_t \mid x_{t-1}^n, \theta)$, yielding $\{\bar{x}_{t-1}^n\}_{n=1}^N$ and *(ii)* propagate these particles to time t by sampling $x_t^n \sim p(x_t \mid \bar{x}_{t-1}^n, y_t, \theta)$ for $n = 1, \ldots, N$.

Finally, by (36), the estimate of the likelihood is given by

$$z := \prod_{t=1}^{T} \left[\frac{1}{N} \sum_{n=1}^{N} \nu_t^n \right]. \tag{38}$$

As for the bootstrap particle filter, it holds that this estimator is non-negative and unbiased. However, it is often the case in practice that (38) has much lower variance than its bootstrap particle filter counterpart (31).

The particle filter described above is often referred to as the *fully adapted auxiliary particle filter*. The term "adapted" here refers to the fact that we adapt the sampling distributions, both in the resampling step and when simulating new particles, to the information provided by y_t. "Fully" refers to the fact that we do so by using the exact conditional distribution in (34). As mentioned above, however, this conditional distribution is only available on closed form for a restricted class of models. More generally it is possible to use an approximation $q(x_t \mid x_{t-1}, y_t, \theta) \approx p(x_t \mid x_{t-1}, y_t, \theta)$ for simulating new particles, and similarly an approximation $q(y_t \mid x_{t-1}, \theta) \approx p(y_t \mid x_{t-1}, \theta)$ when computing the resampling weights. It turns out that it is still possible to obtain a valid particle filter implementation, yielding unbiased estimates of the likelihood, where the approximations q are used as *proposal distributions* within an importance sampling framework. We refer to Pitt and Shephard (1999); Doucet and Johansen (2011) for details.

6.3 Rao-Blackwellization

One important (and rather common) class of state space models for which variance reduction is possible are the so-called *conditionally linear Gaussian (CLG)* state space models. These models are characterised by having a substructure that can be identified as being linear and Gaussian. Therefore, this substructure is analytically tractable using Kalman filtering techniques, which can be leveraged when running a particle filter for the whole model. This is done in a way which resembles the method of *Rao-Blackwellization* of statistical estimators, and the resulting particle filters are therefore commonly referred to as Rao-Blackwellized particle filters (RBPFs).

To be more specific, let the state variable x_t be partitioned into two components $x_t = (x_t^{\mathrm{n}}, x_t^{\mathrm{l}})$ where we use the superscripts n and l for "nonlinear" and "linear", respectively. Then, a CLG model can be defined as a state space model where the *conditional stochastic process* $\{(x_t^{\mathrm{l}}, y_t) \mid x_{0:t}^{\mathrm{n}}\}_{t \geq 0}$ follows a (time-inhomogeneous) linear Gaussian state space model. We can thus identify whether or not a specific model under study is CLG by "pretending" that the state component $\{x_t^{\mathrm{n}}\}_{t \geq 0}$ is observed. If the model then can be viewed as a linear Gaussian state space model, then the original model is CLG.

To give an example, consider the mixed linear/nonlinear model

$$x^{\mathrm{n}}_{t+1} = f^{\mathrm{n}}(x^{\mathrm{n}}_t) + A^{\mathrm{n}}(x^{\mathrm{n}}_t)x^{\mathrm{l}}_t + v^{\mathrm{n}}_t, \tag{39a}$$

$$x^{\mathrm{l}}_{t+1} = f^{\mathrm{l}}(x^{\mathrm{n}}_t) + A^{\mathrm{l}}(x^{\mathrm{n}}_t)x^{\mathrm{l}}_t + v^{\mathrm{l}}_t, \tag{39b}$$

$$y_t = g(x^{\mathrm{n}}_t) + C(x^{\mathrm{n}}_t)x^{\mathrm{l}}_t + e_t \tag{39c}$$

where $x^{\mathrm{n}}_t \in \mathbb{R}^{d_{\mathrm{n}}}$, $x^{\mathrm{l}}_t \in \mathbb{R}^{d_{\mathrm{l}}}$, and $y_t \in \mathbb{R}^{d_y}$. The model is specified in terms of the (possibly nonlinear) functions $f^{\mathrm{n}} : \mathbb{R}^{d_{\mathrm{n}}} \mapsto \mathbb{R}^{d_{\mathrm{n}}}$, $A^{\mathrm{n}} : \mathbb{R}^{d_{\mathrm{n}}} \mapsto \mathbb{R}^{d_{\mathrm{n}} \times d_{\mathrm{l}}}$, $f^{\mathrm{l}} : \mathbb{R}^{d_{\mathrm{n}}} \mapsto \mathbb{R}^{d_{\mathrm{l}}}$, $A^{\mathrm{l}} : \mathbb{R}^{d_{\mathrm{n}}} \mapsto \mathbb{R}^{d_{\mathrm{l}} \times d_{\mathrm{l}}}$, $g : \mathbb{R}^{d_{\mathrm{n}}} \mapsto \mathbb{R}^{d_y}$, and $C : \mathbb{R}^{d_{\mathrm{n}}} \mapsto \mathbb{R}^{d_y \times d_{\mathrm{l}}}$ (where we have oppressed the dependence on the model parameter θ for brevity). Furthermore, the process noises v^{n}_t and v^{l}_t, as well as the measurement noise e_t are all assumed to be Gaussian. To see that this model is indeed CLG, pretend that $\{x^{\mathrm{n}}_t\}_{t \geq 0}$ is observed. Then, all the functions just listed (which specify the model) are known constants and the model is reduced to a linear Gaussian state space model. Note that the dynamical equation (39a) is viewed as a measurement equation in this "pretended" linear Gaussian state space model.

The RBPF exploits the structure of a CLG model by simulating particles representing the "nonlinear state" x^{n}_t, while simultaneously tracking the "linear state" x^{l}_t using a separate Kalman filter for each particle.[5] We do not go in to the details on the RBPF implementation here, but instead refer the interested reader to Chen and Liu (2000); Doucet et al. (2000). Note that CLG models can come in many different forms (hence the rather abstract definition above). However, the variance reduction offered by Rao-Blackwellization can in many cases be substantial, so it is well worth the effort to investigate whether or not a given model under study can be viewed as CLG—thus opening up for using an RBPF—before running a bootstrap particle filter on this model.

7 Numerical illustrations: Learning a nonlinear spring-damper system

In this section, we will walk through a basic, but hopefully illustrative example of probabilistic learning in an applied situation. The code used in the example is available in the appendix and also via the first authors homepage. For a more challenging numerical example, we refer to, e.g., Svensson et al. (2017). The nonlinear spring-damper system that is used to exemplify the method is illustrated in Figure 2 and it is modeled via a spring force F_{s} and a damper force F_{d} according to

$$F_{\mathrm{s}} = -ks^p, \tag{40a}$$

$$F_{\mathrm{d}} = -f_{\mathrm{c}}\mathrm{sign}(\dot{s}) - c_0\dot{s}. \tag{40b}$$

The spring will be linear for $p = 1$, while $p < 1$ results in a nonlinear spring. Furthermore, $f_{\mathrm{c}} = 0$ gives a linear damper, whereas the nonlinear damper $f_{\mathrm{c}} > 0$ is motivated in Spencer et al. (1997). Via the use of force balance

[5] To be more precise, the RBPF maintains particles representing the entire history of the "nonlinear state", $x^{\mathrm{n}}_{0:t}$, and conditionally on these histories the model is linear subject to Gaussian noise.

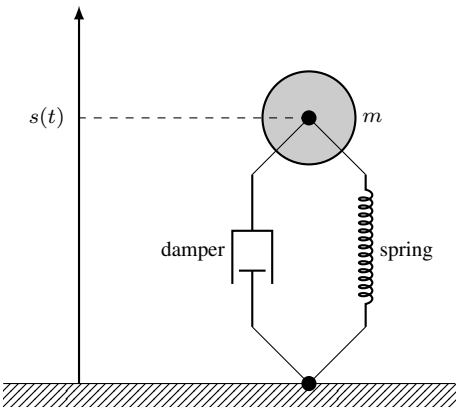

Figure 2: The spring-damper system.

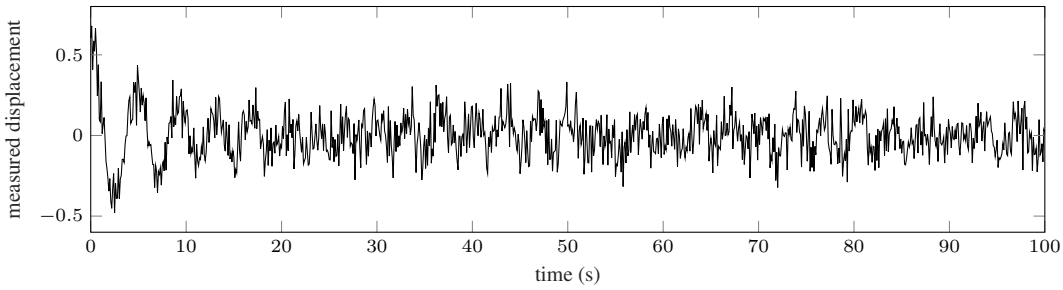

Figure 3: Data simulated from the spring-damper system, which will be used to learn the parameter values θ. As clearly can be seen, there is a significant amount of stochastic noise present in the system, and the data record is not very large.

$\sum F = m\ddot{s}$ and a forward-Euler discretization with sampling time T_s we obtain the following discrete-time state space model

$$x^1_{t+1} = x^1_t + T_s x^2_t, \tag{41a}$$
$$x^2_{t+1} = x^2_t + \tfrac{T_s}{m}(-f_c \text{sign}(x^2_t) - c_0 x^2_t - k(x^1_t)^p) + v_t, \tag{41b}$$
$$y_t = x^1_t + e_t, \tag{41c}$$

where x^1_t represents $s(t)$ and x^2_t represents $\dot{s}(t)$, and we have also included a measurement noise e_t and a process noise v_t (e.g., an external unknown force). From this model, $T = 1\,000$ data points were simulated with an initial known mass displacement of 0.5, shown in Figure 3. Furthermore, the sampling time was $T_s = 0.1$, the mass $m = 8$, and the noise distributions $e_t \sim \mathcal{N}(0, 0.1)$ and $v_t \sim \mathcal{N}(0, 0.2)$, which we all assume is known, whereas the remaining parameters (k, p, f_c, c_0) are assumed unknown. (Their true values used in the simulation are shown in Figure 4 for reference.)

Let us now specify the priors to be used. The parameters f_c, c_0 and k all have to be non-negative by construction, implying that any reasonable prior must be non-negative as well. The Gamma distribution—which we will denote by $\mathcal{G}$—has this property and will therefore be used. More specifically, since f_c is likely to be small we choose $f_c \sim \mathcal{G}(2, 0.01)$, c_0 on the other hand can take on slightly larger values, so its prior is chosen as $\mathcal{G}(2, 1)$. For the spring coefficient k the prior is chosen as $\mathcal{G}(4, .3)$. Finally, p is by physical laws restricted to the interval $0 \leq p \leq 1$, which motivates the prior $\mathcal{U}(0, 1)$. These priors for the parameters are shown in Figure 4 together with their true values. When working with generative models of this kind it a good idea to sample artificial parameter values from the prior and simulate artificial measurements using these parameters in order to see how the model behaves and better understand the prior assumptions made.

With the model definition $p(\theta, x_{0:T}, y_{1:T})$ in place we can now start to learn θ. We did this by generating $10\,000$ samples from $p(\theta \mid y_{1:T})$ using PMH as described in Algorithm 3. The result is shown in Figure 4 (together with the true values and the priors). As proposal in PMH, we used a random walk based on a Gaussian distribution with rather small variance. This takes a few minutes on an ordinary desktop computer. As we have mentioned the PMH algorithm will also provide samples from the state smoothing distribution $p(x_{0:T} \mid y_{1:T})$, which we show in Figure 5.

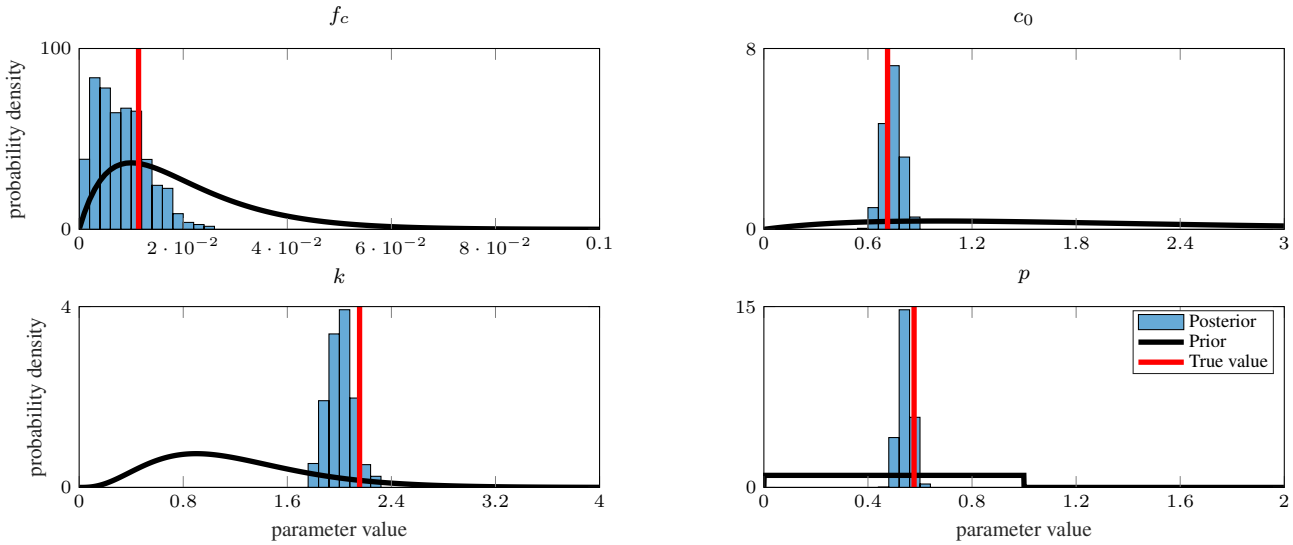

Figure 4: The marginals of the posterior distribution for the parameters, i.e., what can be said about the parameters given data $y_{1:T}$ in Figure 3, the priors (black lines) and the model (40a).

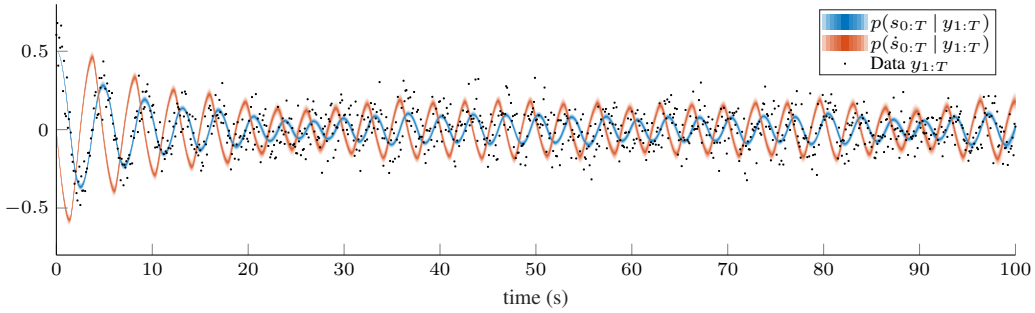

Figure 5: The recorded data (dots) alongside with the computed state smoothing distributions, i.e., $p(x_{0:T} \mid y_{1:T})$ (shaded blue) and $p(\dot{s}_{0:T} \mid y_{1:T})$ (shaded orange).

8 Discussion

Probabilistic modeling is about uncertain representations of data and knowledge using probability distributions and how to actually compute these representation by inference and learning algorithms. We have in this tutorial explained a solution to the problem of computing the posterior distribution of the unknown parameters $p(\theta \mid y_{1:T})$ in a nonlinear state space model (1) using approximate solutions that converge to the truth as we make use of more computational power. This solution involved the use of a particle filter inside a standard Metropolis Hastings algorithm, where the particle filter was used to compute unbiased estimates of the likelihood which in turn enabled the computation of the relevant acceptance probabilities. There are fairly general software implementations available via the modeling language LibBi Murray (2015) which was used to implement the examples. We will end this tutorial by briefly reflecting upon some recent developments on probabilistic *representations* and how to carry out the *computations* required to learn and make use of these representations.

The search and development of new probabilistic models describing nonlinear dynamical phenomena is a highly active and exciting area of research right now. One of the key lessons from modern machine learning is that flexible models often provide the best results Ghahramani (2015). The two dominating approaches for creating flexible models are Bayesian non-parametrics Ghahramani (2013) and deep learning LeCun et al. (2015), both of which can be combined with the state space model to create flexible and highly useful models. The most popular Bayesian non-parametric model is arguable the Gaussian process Rasmussen and Williams (2006) which offers a probabilistic distribution over

functions. The Gaussian process construction thus offers a rather natural probabilistic model for the nonlinear functions $f(\cdot)$ and $g(\cdot)$ in (1). Such a construction was developed in Frigola et al. (2013) and the required learning algorithms were recently enhanced in Svensson and Schön (2017) based on new model approximations. When using flexible models it is important to have a principled way of trading-off their flexibility and the actual fit to data. An emerging way of developing flexible models is provided by probabilistic programming where probabilistic models are represented using computer programs, which was in fact use to implement and solve the examples we gave in this tutorial.

The development of new and more capable approximations based on sequential Monte Carlo methods is progressing fast. As a first clear trend we mention algorithms scaling to higher-dimensional models, where many of the standard algorithms struggle. Relevant developments in this direction involve the so-called block particle filter Rebeschini and van Handel (2015), the location particle smoother Briggs et al. (2013), and various methods reviewed in Djuric and Bugallo (2013). Furthermore, the nested SMC method Naesseth et al. (2015) allows us to *exactly approximate* the locally optimal proposal, and significantly extend the class of models for which we can perform efficient inference using SMC. This development also opens up for approximate inference in nonlinear spatio-temporal state space models.

A probabilistic graphical model Jordan (2004) is a probabilistic representation where a graph encodes and makes use of the underlying structure between the random variables making up the model. Hence, the state space model is a very specific instance of a probabilistic graphical model corresponding to a chain graph, which leads to the rather natural question of whether we can make use of SMC for inference and learning also in more general graphical models. The answer to this question is yes and the key is to introduce a sequential decomposition of the graphical model. This sequential decomposition can then be exploited by SMC, since SMC is in fact applicable to any inference problem that is defined on a sequence of probability distributions defined on a sequence of spaces of increasing dimension. See e.g., Naesseth et al. (2014); Beskos et al. (2014) for developments in this direction. This ends our discussion on the second trend, namely the use of SMC type algorithms for inference in more general probabilistic models.

The third trend is the composition of two or more existing algorithms into new and more capable algorithms. We have in this tutorial portrayed particle Metropolis Hastings algorithm—which is one instantiation of this trend—corresponding to a particular combination of the particle filter and the Metropolis Hastings algorithm. It belongs to the growing family of exact approximation algorithms. The particle Gibbs sampler Andrieu et al. (2010) makes use of the so-called conditional particle filter to generate samples from the posterior distribution. The SMC^2 sampler which was independently derived by Chopin et al. (2013); Fulop and Li (2013) is similar in spirit to the particle Metropolis-Hastings algorithm in that it makes use of a particle filter within another sampling algorithm. More specifically it makes use of an SMC sampler in the θ-dimension that in turn has many internal particle filters in the x-dimension. The particle filter can also be composed with another particle filter to produce interesting and capable algorithms. The idea of coupling two particle filters is developed in Jacob et al. (2017) and as a result is can be used to build a competitive solution to the nonlinear state smoothing problems.

Acknowledgement

This research is financially supported by the Swedish Research Council via the projects *Probabilistic modeling of dynamical systems* (contract number: 621-2013-5524) and *Learning of Large-Scale Probabilistic Dynamical Models* (contract number: 2016-04278), and the Swedish Foundation for Strategic Research (SSF) via the project *ASSEMBLE*.

References

C. Andrieu and G. O. Roberts. The pseudo-marginal approach for efficient Monte Carlo computations. *The Annals of Statistics*, 37(2):697–725, 2009.

C. Andrieu and M. Vihola. Convergence properties of pseudo-marginal Markov chain Monte Carlo algorithms. *The Annals of Applied Probability*, 25(2):1030–1077, 2015.

C. Andrieu, A. Doucet, and R. Holenstein. Particle Markov chain Monte Carlo methods. *Journal of the Royal Statistical Society. Series B (Methodological)*, 72(2):1–33, 2010.

M. A. Beaumont. Estimation of population growth or decline in genetically monitored populations. *Genetics*, 164(3): 1139–1160, 2003.

A. Beskos, D. Crisan, A. Jasra, K. Kamatani, and Y. Zhou. A stable particle filter in high-dimensions. Technical report, ArXiv:1412.3501, 2014.

J. Briggs, M. Dowd, and R. Meyer. Data assimilation for large-scale spatio-temporal systems using a location particle smoother. *Environmetrics*, 24(2):81–97, 2013.

O. Cappé, S. Godsill, and E. Moulines. An overview of existing methods and recent advances in sequential Monte Carlo. *Proceedings of the IEEE*, 95(5):899–924, 2007.

R. Chen and J. S. Liu. Mixture Kalman filters. *Journal of the Royal Statistical Society*, 62(3):493–508, 2000.

N. Chopin, P. E. Jacob, and O. Papaspiliopoulos. SMC^2: an efficient algorithms for sequential analysis of state-space models. *Journal of the Royal Statistical Society, Series B*, 75(3):397–426, June 2013.

J. Dahlin, F. Lindsten, and T. B. Schön. Particle Metropolis Hastings using gradient and Hessian information. *Statistics and Computing*, 25(1):81–92, January 2015.

P. Del Moral. *Feynman-Kac formulae: Genealogical and Interacting Particle Systems with Applications*. Probability and Applications. Springer, New York, USA, 2004.

G. Deligiannidis, A. Doucet, and M. K. Pitt. The correlated pseudo-marginal method. Technical report, arXiv preprint arXiv:1511.04992, 2015.

P. M. Djuric and M. F. Bugallo. Particle filtering for high-dimensional systems. In *IEEE 5th International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP)*, St. Martin, France, 2013.

R. Douc, O. Cappé, and E. Moulines. Comparison of resampling schemes for particle filtering. In *Proceedings of the 4th International Symposium on Image and Signal Processing and Analysis*, 2005.

A. Doucet and A. Johansen. A tutorial on particle filtering and smoothing: Fifteen years later. In D. Crisan and B. Rozovskii, editors, *The Oxford Handbook of Nonlinear Filtering*, pages 656–704. Oxford University Press, Oxford, UK, 2011.

A. Doucet, S. J. Godsill, and C. Andrieu. On sequential Monte Carlo sampling methods for Bayesian filtering. *Statistics and Computing*, 10(3):197–208, 2000.

R. Frigola, F. Lindsten, T. B. Schön, and C. E. Rasmussen. Bayesian inference and learning in Gaussian process state-space models with particle MCMC. In *Advances in Neural Information Processing Systems (NIPS) 26*, Lake Tahoe, NV, USA, December 2013.

A. Fulop and J. Li. Efficient learning via simulation: A marginalized resample-move approach. *Journal of Econometrics*, 176(2):146–161, October 2013.

A. Gelman, J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. B. Rubin. *Bayesian data analysis*. CRC Press, third edition, 2013.

Z. Ghahramani. Bayesian non-parametrics and the probabilistic approach to modelling. *Philosophical Transactions of The Royal Society A Mathematical Physical and Engineering*, 371(20110553), 2013.

Z. Ghahramani. Probabilistic machine learning and artificial intelligence. *Nature*, 521(7553):452–459, May 2015.

N. J. Gordon, D. J. Salmond, and A. F. M. Smith. Novel approach to nonlinear/non-Gaussian Bayesian state estimation. In *IEE Proceedings on Radar and Signal Processing*, volume 140, pages 107–113, 1993.

W. K. Hastings. Monte Carlo simulation methods using Markov Chains and their applications. *Biometrica*, 57:97–109, 1970.

J. D. Hol, T. B. Schön, and F. Gustafsson. On resampling algorithms for particle filters. In *Proceedings of the Nonlinear Statistical Signal Processing Workshop*, Cambridge, UK, September 2006.

P. E. Jacob, F. Lindsten, and T. B. Schön. Smoothing with couplings of conditional particle filters. Technical report, arXiv:1701.02002, 2017.

M. I. Jordan. Graphical models. *Statistical Science*, 19(1):140–155, 2004.

N. Kantas, A. Doucet, S. S. Singh, J. M. Maciejowski, and N. Chopin. On particle methods for parameter estimation in state-space models. *Statistical Science*, 30(3):328–351, 2015.

G. Kitagawa. Monte Carlo filter and smoother for non-Gaussian nonlinear state space models. *Journal of Computational and Graphical Statistics*, 5(1):1–25, 1996.

Y. LeCun, Y. Bengio, and H. G. Deep learning. *Nature*, 521:436–444, 2015.

N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller. Equations of state calculations by fast computing machine. *Journal of Chemical Physics*, 21(6):1087–1092, 1953.

S. P. Meyn and R. L. Tweedie. *Markov chains and stochastic stability*. Cambridge University Press, 2009.

L. M. Murray. Bayesian state-space modelling on high-performance hardware using LibBi. *Journal of Statistical Software*, 67(10):1–36, 2015.

A. C. Naesseth, F. Lindsten, and T. B. Schön. Sequential Monte Carlo for graphical models. In *Advances in Neural Information Processing Systems (NIPS) 27*, Montreal, Quebec, Canada, December 2014.

A. C. Naesseth, F. Lindsten, and T. B. Schön. Nested sequential Monte Carlo methods. In *Proceedings of the 32nd International Conference on Machine Learning (ICML)*, Lille, France, July 2015.

A. B. Owen. *Monte Carlo theory, methods and examples*. 2013. URL `http://statweb.stanford.edu/~owen/mc/`. Book draft.

V. Peterka. Bayesian system identification. *Automatica*, 17(1):41–53, 1981.

M. K. Pitt and N. Shephard. Filtering via simulation: Auxiliary particle filters. *Journal of the American Statistical Association*, 94(446):590–599, 1999.

M. K. Pitt, R. dos Santos Silva, R. Giordani, and R. Kohn. On some properties of Markov chain Monte Carlo simulation methods based on the particle filter. *Journal of Econometrics*, 171(2):134–151, 2012.

C. E. Rasmussen and C. K. I. Williams. *Gaussian processes for machine learning*. MIT Press, 2006.

P. Rebeschini and R. van Handel. Can local particle filters beat the curse of dimensionality? *Annals of Applied Probability*, 25(5):2809–2866, 2015.

C. Robert and G. Casella. *Monte Carlo Statistical Methods*. Springer-Verlag New York, 2004.

C. Robert and G. Casella. A short history of Markov chain Monte Carlo: subjective recollections from incomplete data. *Statistical Science*, 26(1):102–115, 2011.

T. B. Schön, F. Lindsten, J. Dahlin, J. Wågberg, A. C. Naesseth, A. Svensson, and L. Dai. Sequential Monte Carlo methods for system identification. In *Proceedings of the 17th IFAC Symposium on System Identification (SYSID)*, Beijing, China, October 2015.

B. F. Spencer, S. J. Dyke, M. K. Sain, and J. D. Carlson. Phenomenological model for magnetorheological dampers. *Journal of engineering mechanics*, 123(3), 1997.

L. Stewart and P. McCarty. The use of Bayesian belief networks to fuse continuous and discrete information for target recognition and discrete information for target recognition, tracking, and situation assessment. In *Proceedings of SPIE Signal Processing, Sensor Fusion and Target Recognition*, volume 1699, pages 177–185, 1992.

A. Svensson and T. B. Schön. A flexible state space model for learning nonlinear dynamical systems. *Automatica*, 2017. (Accepted for publication).

A. Svensson, F. Lindsten, and T. B. Schön. Learning nonlinear state-space models with highly informative observations: a tempered sequential Monte Carlo solution. *submitted to Mechanical Systems and Signal Processing (MSSP)*, 2017.

A Implementation in LibBi

The numerical example was implemented using LibBi (Murray (2015), `www.libbi.org`). The code needed to reproduce the examples will be available via the first authors homepage. Apart from simulating data, the LibBi code for implementing the model (40a) and run the PMH learning algorithm is provided below.

First, the following content is placed in a file called `model.bi`.

```
model Damper {
  /* Specifying all variables in the problem */
  const T_s = .1, m = 2          /* Known constants */
  param f_c, c_0, k, p           /* Unknown variables */
  noise v                        /* Process noise */
  state s, sdot, s_old, sdot_old  /* State variables */
  obs y

  sub parameter {
    /* Specifying all priors for the unknown parameters  */
    f_c ~ gamma(2, 0.01)
    c_0 ~ gamma(2, 1)
    k ~ gamma(4, 0.3)
    p ~ uniform(0, 1)
  }

  sub proposal_parameter {
    /* Specifying the random walk proposal used in PMH */
    f_c ~ truncated_gaussian(f_c, 1.0e-3, 0.0)
    c_0 ~ truncated_gaussian(c_0, 1.0e-2, 0.0, 1.0)
    k ~ truncated_gaussian(k, 1.0e-2, 0.0)
    p ~ truncated_gaussian(p, 1.0e-2, 0.0, 1.0)
  }

  sub initial {
    /* Specifying the inital state variables */
    s <- 0.5
    sdot <- 0.0
  }

  sub transition(delta = 1) {
    /* Describe how the states in the model evolves at each time step. The variables s_old and sdot_old is only for
         temprary storage of the previous state values. */
    v ~ gaussian(0.0, 1.0e-2)

    s_old <- s
    sdot_old <- sdot

    s <- s_old + T_s*(sdot_old)
    sdot <- sdot_old + (T_s/m)*(-f_c*((sdot_old < 0) ? -1 : 1) - c_0*sdot_old - k*((s_old < 0) ? -1 : 1)*pow(abs(s_old
         ), p)) + v
  }

  sub observation {
    /* Describe how the measurements relates to the states */
    y ~ gaussian(s, 1.0e-1)
  }
}
```

Once the file `model.bi` is in place, the following command can be run from the command line:

```
libbi sample --target posterior --model-file Damper.bi --nsamples 10000 --nparticles 256 --end-time 1000 --sampler mh
     --obs-file data/obs.nc --output-file results/posterior.nc
```

which will produce 10 000 samples from $p(\theta \mid y_{1:1000})$, where $y_{1:1000}$ is found in the file `data/obs.nc`. The samples are generated by the PMH algorithm and recorded in the file `results/posterior.nc` which can be loaded into, e.g., Matlab or R for further processing.